

UTILITY APPLICATION  
of  
Melvin Richard Zimowski  
for  
UNITED STATES LETTERS PATENT  
on  
TECHNIQUE FOR MAINTAINING AND MANAGING  
DYNAMIC WEB PAGES STORED IN A SYSTEM CACHE  
AND REFERENCED OBJECTS CACHED IN OTHER DATA STORES

Docket No. P70 43480  
Drawings: 2

Attorneys  
PRETTY, SCHROEDER & POPLAWSKI  
444 South Flower Street, 19th Floor  
Los Angeles, California 90071  
Ofc: 213/622-7700  
Fax: 213/489-4210

**CERTIFICATE OF MAILING BY "EXPRESS MAIL"**

"EXPRESS MAIL" MAILING LABEL NUMBER EL 243841515 US  
DATE OF DEPOSIT June 23, 2000

I HEREBY CERTIFY THAT THIS PAPER OR FEE IS BEING DEPOSITED WITH THE UNITED STATES POSTAL SERVICE "EXPRESS MAIL POST OFFICE TO ADDRESSEE" SERVICE UNDER 37 CFR 1-10 ON THE DATE INDICATED ABOVE AND IS ADDRESSED TO BOX PATENT APPLICATION, ASSISTANT COMMISSIONER FOR PATENTS, WASHINGTON, D. C. 20231.

Joanne H. Housen

(TYPED OR PRINTED NAME OR PERSON MAILING PAPER OR FEE)

Joanne H. Hansen

(SIGNATURE OF PERSON MAILING PAPER OR FEE)

# TECHNIQUE FOR MAINTAINING AND MANAGING DYNAMIC WEB PAGES STORED IN A SYSTEM CACHE AND REFERENCED OBJECTS CACHED IN OTHER DATA STORES

## PROVISIONAL APPLICATION

5 This application claims the benefit of U.S. Provisional Application No. 60/140,711 entitled  
"TECHNIQUE FOR MAINTAINING AND MANAGING DYNAMIC WEB PAGES STORED IN  
A SYSTEM CACHE AND REFERENCED OBJECTS CACHED IN OTHER DATA STORES,"  
filed on June 24, 1999, by Mel Zimowski, attorney's reference number ST9-99-080, which is  
incorporated by reference herein.

## BACKGROUND OF THE INVENTION

## 1. FIELD OF THE INVENTION

This invention relates in general to computer-implemented systems, and, in particular, to maintaining and managing dynamic web pages and objects referenced by the web pages.

## 2. DESCRIPTION OF RELATED ART

The Internet computer network is a collection of computer networks that exchange information via the Transmission Control Protocol/Internet Protocol ("TCP/IP") protocol suite. Currently, the use of the Internet computer network for commercial and non-commercial uses is exploding. Via its networks, the Internet computer network enables users in different locations to access information stored in data sources (e.g., databases) on servers distributed across these networks.

The World Wide Web (i.e., the "WWW" or the "Web") is a hypertext information and communication system used on the Internet computer network with data communications operating according to a client/server model. Typically, a user of a Web browser at a Web client computer will request data stored in data sources from a Web server computer, at which Web server software resides. The Web server software interacts with other computer programs that

use interfaces to connect to these data sources, for example, a database managed by a Database Management System ("DBMS"), or uses the interfaces directly to access these data sources. These computer programs residing at the Web server computer transmit the requested data to the client computer in worldwide web documents referred to as web pages. The data can be of many 5 different types of information, including database data, images, video clips, or audio tracks.

Web pages can be static web pages (i.e. web pages with fixed content that are pre-generated long before the Web client request is issued) or dynamic web pages (i.e., web pages whose content is dynamically generated at the time the web client request is processed).

Dynamic web pages are typically expensive to generate because they contain data that must be obtained dynamically at web servers from either local or remote data sources. For this 10 reason, web server caches are often used to store dynamic Web pages that are requested by multiple users. These caches are of finite size and have limits on the number of web pages they can contain. Further, the need for retaining individual pages within the cache varies over time. Thus, without proper maintenance, the storage allocated to the cache may become completely 15 used up or the content of the cache may become outdated.

A web page may and often does contain hypertext links to objects stored in data stores (e.g., a local file system) that can be accessed by the web server that serves the web page. These 20 hypertext links permit the objects that they reference to be made available to a user at a web browser as an integral part of the web page that contains them. Thus, the web page appears incomplete if it is served to a user and the objects referenced through these hypertext links cannot be materialized as required by the web page or user.

When web pages are generated dynamically, it is not possible to predetermine the objects that will be referenced within these web pages using hypertext links. In addition, when dynamic 25 web pages are stored in a system cache for reuse, the relationship between the web page and the objects that it references is no longer momentary. Thus, unless the contents of the system cache and the data stores containing the objects referenced in the web page are properly managed, incomplete web pages could easily be displayed at web browsers. Accordingly, the providers of web sites need to manage the contents of a dynamic web page system cache and the contents of

the data stores containing objects that are referenced by hypertext links that occur within those dynamic web pages in a coordinated fashion.

Thus, there is a need in the art for improved techniques for maintaining and managing dynamic web pages stored in system caches and objects referenced by the web pages.

5

## SUMMARY OF THE INVENTION

To overcome the limitations in the prior art described above, and to overcome other limitations that will become apparent upon reading and understanding the present specification, the present invention discloses a method, apparatus, and article of manufacture for managing data stored in a data storage device connected to a computer.

In accordance with the present invention, it is determined that a web page is to be cached. The web page references other objects. The referenced objects are stored in one or more data stores. The web page is cached. The cached web page and the referenced objects are managed in a coordinated fashion to ensure the display of a complete web page.

## BRIEF DESCRIPTION OF THE DRAWING

Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

FIG. 1 schematically illustrates the hardware environment of a preferred embodiment of the present invention, and more particularly, illustrates a typical distributed computer system using the Internet; and

FIG. 2 is a flow chart illustrating the steps performed in accordance with an embodiment of the present invention.

## DETAILED DESCRIPTION

In the following description, reference is made to the accompanying drawings which form a part hereof, and which is shown by way of illustration a specific embodiment in which the invention may be practiced. It is to be understood that other embodiments may be utilized as structural changes may be made without departing from the scope of the present invention.

Sub A1/15

~~10~~ A. Overview

One embodiment of the present invention provides a management system for maintaining and managing dynamic web pages and objects referenced in the web pages. The preferred embodiment of the invention provides techniques for managing the contents of a dynamic web page system cache and related data stores containing objects that the web pages reference while also ensuring the completeness of cached web pages that are subsequently reused for display at browsers.

The preferred embodiment of the invention uses a DBMS (e.g., DB2 from International Business Machines, Corporation) to cache dynamic web pages and to track dependencies that the dynamic web pages have on large objects cached in the local UNIX file system (in particular, the Hierarchical File System available under OS/390 UNIX System Services). The preferred embodiment of the invention provides tools for managing the contents of the dynamic web page cache and the large object cache, while also ensuring that the management is performed in such a way as to guarantee the completeness of any cached web pages that are displayed at a browser.

Sub A2/20

~~10~~ B. Hardware Environment

FIG. 1 schematically illustrates the hardware environment of a preferred embodiment of the present invention, and more particularly, illustrates a typical distributed computer system using the Internet 100 to connect Web client computers 102 executing Web browsers to a Web server computer 104 executing Web server software and other computer programs that connect the server system 104 to data sources 106. A typical combination of resources may include client computers 102 that are personal computers or workstations, and a web server computer 104 that is a personal computer, workstation, minicomputer, or mainframe. These systems are coupled to one another by various networks, including LANs, WANs, SNA networks, and the Internet.

Sub A3/25

A Web client computer 102 typically executes a Web browser and is coupled to a Web server computer 104 executing Web server software. The Web browser is typically a program such as Microsoft's Internet Explorer or Netscape Navigator. The Web server software is typically a program such as IBM's HTTP Server or other WWW server software. The software

Cont  
A3 > executing on the Web server uses a data source interface and, possibly, other computer programs, for connecting to the data sources 106. The software executing on the Web server may also include a cache management system 110. The client computer 102 is bi-directionally coupled with the Web server computer 104 over a line or via a wireless system. In turn, the Web server computer 104 is bi-directionally coupled with data sources 106.

5 The data source interface permits the software executing on the Web Server to be connected to a Database Management System (DBMS), which supports access to a data source 106 by executing DBMS software. The DBMS may be located on the same server as the Web server computer 104 or may be located on a separate machine. The data sources 106 may be 10 geographically distributed. The software executing on the Web server translates the request received from a Web browser into one or more statements (e.g., a macro file or a COBOL program) that can be processed to retrieve data from data sources 106.

Those skilled in the art will recognize many modifications may be made to this configuration without departing from the scope of the present invention.

15 C. Dynamic Web Page Caching Overview

1. Overview of Design

20 The caching of dynamic web pages enhances the Web server's ability to quickly serve web pages containing dynamic content. In the preferred embodiment, ASCII encoded web pages are cached in a database for reuse by the cache management system 110. This eliminates the costs associated with reconstructing a dynamic web page that is requested by multiple users. Web pages can be HTML web pages or XML documents.

A cached page is static, as the content of a cached page reflects the state of data stores and business logic at the time the web page was created. Subsequent changes to the data stores and business logic do not affect the content of the cached page.

25 Caching directives are used to specify, among other things, the web pages to be cached. A number of factors affect whether web pages should be cached. More specifically, a page should be cached when the page is repeatedly requested by users and when the content of the page changes infrequently. A pages should not be cached when the processing associated with

the generation of the web page makes changes to data sources. If a cached web page is used to respond to a user's request, the processing logic associated with the generation of the cached page is not executed and no changes are made to the data sources.

When the cache management system 110 (also referred to as the management system) is configured as a CGI application, the caching directives are processed for each cache management system user request. When the cache management system 110 is configured as a Web server plugin or as a servlet, the caching directives are processed once per web server address space at the time the first management system request is assigned to a worker thread associated with that address space. For additional information about worker threads, refer to pending U.S. Patent application, Serial No. 09/104,879, filed June 25, 1998 by M. Zimowski, S. Greenspan, P. Livecchi, and J. Aman, entitled "METHOD AND SYSTEM FOR MANAGING CONNECTIONS TO A DATABASE MANAGEMENT SYSTEM".

Stored procedures are provided by the cache management system 110 for managing the contents of the cache.

2. Web Page Caching Directives

In the preferred embodiment, authorized Web server administrators configure the cache management system 110 to cache web pages by adding caching directives to the management system initialization file (db2www.ini).

In particular, the DTW\_CACHE\_PAGE directive is used to specify web pages that are to be cached by the cache management system 110. If the management system initialization file does not contain a DTW\_CACHE\_PAGE directive, then no web pages are cached.

The following is the syntax for the DTW\_CACHE-PAGE directive:

DTW\_CACHE\_PAGE [=] file\_name\_spec|path\_template\_spec lifetime usage\_scope

The term file\_name\_spec refers to the specification of one or all blocks within a macro file using the fully qualified name of the macro file. A macro file is an installation-provided application that the cache management system 110 executes to generate one or more web pages.

A block is a subsection of a macro that is capable of generating a specific web page. For example, DTW\_CACHE\_PAGE = /u/USRND01/macros/custqord.d2w/Output specifies the caching of the web page created by the execution of the block Output in macro custqord.d2w in the directory /u/USRND01/macros. In the following additional example, DTW\_CACHE\_PAGE = /u/USRND01/macros/custqord.d2w/\* specifies the caching of all web pages created by the execution of any block in macro custqord.d2w in the directory /u/USRND01/macros.

The term path\_template\_spec refers to the specification of blocks within macro files using a path template for one or more directories containing macro files. A path template contains the suffix/\*. The management system caches all web pages created by the execution of blocks in macro files contained within the directory or directories that match the path template. For example, DTW\_CACHE\_PAGE = /u/USRND01/macros/\* specifies the caching of web pages created by the execution of blocks in all macros contained within the directory /u/USRND01/macros and any subdirectory of /u/USRND01/macros. In the following additional example, DTW\_CACHE\_PAGE = /\* specifies the caching of all web pages created by all management system macros.

The term lifetime refers to the minimum number of seconds that a cached web page is valid.

The term usage\_scope specifies the degree to which the reuse of the web page is restricted. Reuse is granted or denied based on the authority of the userid associated with the request. Usage\_scope can have a value of PUBLIC or PRIVATE. PUBLIC means that the cached web page should be served (i.e., returned to a user) when the user request matches the cache key (discussed in further detail below), the cached page is valid, and the userid associated with the web server thread or process processing the request is authorized to execute the macro that generated the page. PRIVATE means that the cached web page should be served when the user request matches the cache key (discussed in further detail below), the cached page is valid, and the userid associated with the web server thread or process processing the request is the same as the userid that was associated with the web server thread or process that cached the web page.

The caching directive can be specified multiple times. Namely, a different caching directive (i.e., DTW\_CACHE\_PAGE directive) can be specified for each file\_name\_spec or

path\_template\_spec value. When the caching directives conflict with each other, the first directive specified takes precedence.

A cached page is reused for a request when the URL, the form data, and the query string of the request match the URL, form data, and query string of the request that caused the web page to be cached. Examples of caching directives are shown below:

5 Example 1: Specifies the caching of any Web pages generated when the cache management system executes a particular HTML block in the specified macro

```
DTW_CACHE_PAGE /u/USER1/macros/main.d2w/output 3600 PUBLIC
```

In this example, the cache management system caches the Web pages generated when it executes the output block in the macro main.d2w, located in the /u/USER1/macros directory. The Web pages have PUBLIC scope, and remain valid for at least 1 hour.

Example 2: Specifies the caching of any Web pages generated when the cache management system executes any block in the specified macro

```
DTW_CACHE_PAGE /u/USER1/macros/main.d2w/* 1800 PUBLIC
```

15 In this example, the cache management system caches any Web pages the cache management system generates when it executes any block in the macro main.d2w, located in the /u/USER1/macros directory. The Web pages have PUBLIC scope, and remain valid for at least 30 minutes.

20 Example 3: Specifies the caching of any Web pages generated when the cache management system executes blocks in macros located in one or more directories

```
DTW_CACHE_PAGE /u/USER1/macros/* 3600 PRIVATE
```

In this example, the cache management system caches any Web pages the cache management system generates when it executes any block in any macro located in the /u/USER1/macros directory or any of its subdirectories. The Web pages have PRIVATE scope, and remain valid for at least 1 hour.

5 Example 4: Specifies the caching of any Web page generated by all macros

```
DTW_CACHE_PAGE /* 3600 PUBLIC
```

In this example, the cache management system caches all Web pages that the cache management system generates. The Web pages have PUBLIC scope, and remain valid for at least 1 hour.

10 Example 5: Specifies multiple Web pages caching directives

```
DTW_CACHE_PAGE /u/USER1/macros/main/* 1800 PUBLIC
```

```
DTW_CACHE_PAGE /u/USER1/macros/special/daily_news.d2w/* 43200 PUBLIC
```

```
DTW_CACHE_PAGE /u/USER1/macros/special/employee_stats.d2w/* 3600 PRIVATE
```

In this example, the cache management system caches the following: (1) All Web pages generated from any block in any macro located in the /u/USER1/macros/main/ directory. The Web pages have PUBLIC scope and remain valid for at least 30 minutes. (2) All Web pages generated by the daily\_news.d2w macro in the directory /u/USER1/macros/special/. These Web pages have PUBLIC scope and remain valid for at least 12 hours. (3) All Web pages generated by the employee\_stats.d2w macro in the directory /u/USER1/macros/special/. These Web pages have PRIVATE scope and remain valid for at least 1 hour.

15 After the web pages to be cached have been specified, the table used to cache the web pages must be set up. The following steps are necessary to set up the table used to cache the web

pages:

- a. Create the Web page cache table, SYSIBM.DTWCACHEPAGES.
- b. Define the stored procedure used to insert the cached pages into SYSIBM.DTWCACHEDPAGES to the DBMS.
- 5 c. Prepare the stored procedure for execution using a user ID with INSERT, SELECT, and DELETE privileges on SYSIBM.DTWCACHEDPAGES. The user IDs associated with the requests that cache pages must have the EXECUTE privilege for the stored procedure.

After the above three steps are completed, web pages can be cached.

### 3. Web Page Cache Key

The actual cache key for the cached dynamic web page consists of path information, macro name, HTML or XML block name, plus the query string plus the form data (if present) that caused the dynamic web page to be generated.

## D. Management System Dynamic Web Page Cache and Referenced Object Data Store Consistency

### 1. Management System Objectives Regarding Cache Consistency

A web page may be defined as a set of HTML tags or an XML (Extensible Mark-up Language) document and the objects (e.g., DB2 large objects or LOBs) referenced by the web page using hypertext links.

20 The contents of a cached web page reflect the state of data stores and business logic at the time the cached web page was created. Thus, a cached web page is static.

A cached web page is automatically deleted upon expiration of its lifetime (i.e., lifetime value plus creation time), when automatic management is enabled, or until deleted by an authorized web server administrator.

25 The following are elements that the management system is designed to achieve regarding dynamic web page cache and referenced object consistency. When serving a cached web page,

the management system maintains the consistency that existed at the time the cached web page was created. The removal of a referenced object from its data store (e.g. the local UNIX file system) invalidates the cached dynamic web page that references the object. The removal of a dynamic web page from the management system cache invalidates referenced objects stored in the referenced object data store.

A web page is not returned to a user at a browser until all LOBs referenced in the web page are successfully placed in the referenced object data store. Additionally, the web page is not cached until all LOBs referenced in the web page are successfully placed in the referenced object data store. When a web page is removed from the dynamic Web page cache, any LOBs referenced by that web page are removed from the referenced object data store. When a LOB is removed from the referenced object data store, any web page that references the LOB is removed from the dynamic Web page cache and any other LOBs referenced by that web page are removed from the referenced object data store. The cached web page is always removed before any dependent LOBs are removed from the referenced object data store.

## 2. The Management System Dynamic Web Page Cache

ASCII encoded web pages are cached in a data store for reuse by the management system. In the preferred embodiment, referenced objects are large objects (LOBs). Caching directives specify the web pages to be cached and the minimum number of seconds the cached page is valid. The key of a cached page consists of path information, macro name, HTML or XML block name, plus query string plus form data (if present). Cached web pages are classified by usage scope: PUBLIC or PRIVATE.

The DB2 dynamic web page cache table is: SYSIBM.DTWCACHEDPAGES. The index key for a cached page in the dynamic web page cache table is: INDEXED\_KEY CHAR(250), which is the first 250 characters of path information, macro name, HTML or XML block name, query string, and form data. An identifier is specified by: ID INTEGER, which is an identifier derived from path information, macro name, HTML or XML block name, query string, and form data. To provide fast access to the cache, an index to the dynamic web page cache table is provided that consists of both the column for the INDEXED\_KEY and the column for the ID.

The actual key of a cached page is: ACTUAL\_KEY VARCHAR(4000), which is comprised of the path information, macro name, HTML or XML block name, query string, and form data for the request that created the cached page.

The userid of the creator is: CREATOR CHAR(8), which is the userid associated with a request that created the cached page. The creation timestamp is: CREATION\_TIME 5 TIMESTAMP, which specifies the date and time of the creation of the cached page. This date and time is the same as the CREATION\_TIME for any LOBs that the web page references.

The expiration timestamp is: EXPIRATION\_TIME TIMESTAMP, which specifies the date and time of expiration of the cached page (value of CREATION\_TIME + lifetime value 10 from DTW\_CACHE\_PAGE directive). A second index for the dynamic web page cache table is provided that consists of the column for EXPIRATION\_TIME. The indexed column (i.e., EXPIRATION\_TIME) of the second index is used to efficiently identify the cached web pages 15 that have expired.

The size is: SIZE INTEGER, which specifies the size of the cached page in bytes. The usage scope is: USAGE\_SCOPE SMALLINT; a value of 1 means that the page has a PUBLIC 20 usage scope and a value of 2 means that the page has a PRIVATE usage scope. The ordinal position of the segment is: ORDINAL\_POSITION SMALLINT, which specifies the ordinal position of the web page segment within the complete cached page. The dynamic web page segment is: PAGE\_SEGMENT VARCHAR(28100) FOR BIT DATA, which is the ASCII encoded web page segment.

### 3. Dynamic Web Page/LOB Dependency Table

A dynamic web page/LOB dependency table must be created and several stored procedures must be installed before the cache management system 110 can manage cached web pages and LOBs. The steps necessary to create the web page dependency table are outlined 25 below:

a. The dynamic Web page/LOB dependency table contains information about the LOB files stored in HFS and about the relationship that these files may have, if any, to Web pages stored in the dynamic Web page cache.

(i) Create the dynamic Web Page/LOB dependency table, SYSIBM.DTWCACHEDEPS.

(ii) Define the stored procedure used to insert the dependency information into SYSIBM.DTWCACHEDEPS to the DBMS.

5 (iii). Prepare the stored procedure for execution using a user ID with INSERT, SELECT, and DELETE privileges on SYSIBM.DTWCACHEDEPS. The user IDs associated with the requests that retrieve LOBs must have the EXECUTE privilege for the stored procedure.

b. Create the stored procedure that performs automatic management of the Web page cache and LOBs.

10 (i) Define the stored procedure to the DBMS.

(ii) Prepare the stored procedure for execution using a user ID with the DELETE privilege on SYSIBM.DTWCACHEPAGES and SYSIBM.DTWCACHEDEPS. The user IDs that execute macros must have the EXECUTE privilege for the stored procedure.

15 c. Create the stored procedure that is used for more advanced management of the Web page cache and LOBs.

(i) Define the stored procedure to the DBMS.

(ii) Prepare the stored procedure for execution using a userid with the DELETE privilege on SYSIBM.DTWCACHEPAGES and on SYSIBM.DTWCACHEDEPS. The userids that execute the management system-provided macro must have the EXECUTE privilege for the stored procedure.

20

#### 4. Dynamic Web Page/LOB Dependency Table Details

The DB2 web page/LOB dependency table is: SYSIBM.DTWCACHEDEPS. The indexed key for a cached page in the DB2 web page/LOB dependency table is: INDEXED\_KEY CHAR(250), which specifies the first 250 characters of path information, macro name, HTML or XML block name, query string, and form data for the request that created the cached page. The identifier is: ID INTEGER, which specifies an identifier derived from path information,

macro name, HTML or XML block name, query string, and form data for the request that created the cached page. An index to the DB2 web page/LOB dependency table is provided that consists of both the column for the INDEXED\_KEY and the column for the ID.

5 Actual key of the cached page is: ACTUAL\_KEY VARCHAR(4000) - path information, macro name, HTML or XML block name, query string, and form data for the request that created the cached page. The fully qualified HFS filename for the LOB is: FILENAME VARCHAR(1024).

10 The creation timestamp is: CREATION\_TIME TIMESTAMP, which specifies the date and time of the creation of the LOB. This date and time is the same as the CREATION\_TIME for the web page that references the LOB.

15 The expiration timestamp is: EXPIRATION\_TIME TIMESTAMP, which specifies the date and time for expiration of the LOB (value of CREATION\_TIME + DTW\_LOB\_LIFETIME configuration value, when dynamic web page caching not in use; value of CREATION\_TIME + max(DTW\_LOB\_LIFETIME configuration value, lifetime, value from DTW\_CACHE\_PAGE directive) when dynamic web page caching in use). A second index for the DB2 web page/LOB dependency table is provided that consists of the column for EXPIRATION\_TIME. The indexed column (i.e., EXPIRATION\_TIME) of the second index is used to efficiently identify objects (referenced by a web page) that have expired.

20 The size is: SIZE INTEGER, which specifies the size of the LOB in bytes.

The SYSIBM.DTWCACHEDPAGES and the SYSIBMDTWCACHEDEPS DB2 tables are shown below. SYSIBM.DTWCACHEDPAGES is the dynamic Web page cache table. This table contains the cached Web pages and information about the cached Web pages.

Table 1. SYSIBM.DTWCACHEDPAGES

Column: Column Datatype	Description
INDEXED_KEY: CHAR(250)	Indexed key for cached page: the first 250 characters of the actual key
5 ID: INTEGER	Identifier: an identifier derived from the key
ACTUAL_KEY: VARCHAR(4000)	The actual key of the cached page: the path information, macro name, HTML or XML block name, query string, and form data of the request that generated the page.
CREATOR: CHAR(8)	User ID of creator: user ID associated with the request that created the cached page.
CREATION_TIME: TIMESTAMP	Creation timestamp: date and time of creation of cached page.
EXPIRATION_TIME: TIMESTAMP	Expiration timestamp: date and time for expiration of cached page (value of CREATION_TIME + lifetime value from DTW_CACHE_PAGE directive)
SIZE: INTEGER	Size: size of cached page in bytes
15 USAGE_SCOPE: SMALLINT	Usage scope: a value of 1 means that the page has a PUBLIC usage scope and a value of 2 means that the page has a PRIVATE usage scope
ORDINAL_POSITION: SMALLINT	Ordinal position of segment: the ordinal position of the Web page segment within the complete cached page
PAGE_SEGMENT VARCHAR(28100)FOR BIT DATA	Dynamic Web page segment: the ASCII encoded Web page segment
20	

SYSIBM.DTWCACHEDEPS is the Web page dependency table. This table contains information about the LOBs referenced by Web pages.

Table 2. SYSIBM.DTWCACHEDEPS

Column: Column Datatype	Description
5 INDEXED_KEY: CHAR(250)	Indexed key for Web page: first 250 characters of the actual key
ID: INTEGER	Identifier: an identifier derived from the actual key
ACTUAL_KEY: VARCHAR(4000)	The actual key of the cached page: the path information, macro name, HTML or XML block name, query string, and form data of the request that generated the page.
FILENAME: VARCHAR(1024)	Fully qualified HFS filename for the LOB
CREATION_TIME:  TIMESTAMP	Creation timestamp: date and time of the creation of the LOB.  Is the same as the CREATION_TIME for the Web page that references this LOB
EXPIRATION_TIME:  15 TIMESTAMP	Expiration timestamp: date and time of the expiration of the LOB (value of CREATION_TIME + DTW_LOB_LIFETIME configuration value when dynamic web page caching not in use; value of CREATION_TIME + max( DTW_LOB_LIFETIME configuration value, lifetime value from DTW_CACHE_PAGE directive) when dynamic web page caching in use)
SIZE: INTEGER	Size: size of the LOB in bytes

## 5. Managing LOBs and Cached Web Pages

LOBs referenced in web pages are stored as files in a single HFS directory. The name of the directory is specified by the HTML\_PATH configuration statement of the management system initialization file.

5 The settings of two configuration variables control the manner in which LOBs and cached web pages are automatically managed. The configuration variable DTW\_LOB\_LIFETIME specifies the minimum number of seconds LOBs are available. The default is 0 seconds. When dynamic web page caching is in use, the minimum number of 10 seconds a LOB is available is the larger of the value of DTW\_LOB\_LIFETIME and the lifetime value specified in the caching directive for the dynamic web page that references the LOB. The configuration variable DTW\_CACHE\_MANAGEMENT\_INTERVAL specifies a minimum number of seconds between successive automatic cache management attempts. The default is 0 seconds. A value of 0 means that the cache management system 110 will not perform automatic cache management.

15 Cached web pages and LOBs are automatically managed by the management system using DB2 stored procedures. Additionally, a management system-provided macro permits Web server administrators to manage LOBs and cached web pages in more sophisticated ways 20 using a stored procedure.

## 6. Automatic Cache Management

When automatic cache management occurs, all cached dynamic web pages and LOBs with expiration timestamps that precede the current timestamp are automatically deleted.

The following are the logical steps for automatic cache management:

- # Delete cache table entries where expiration timestamp < current timestamp
- # Select dependency table entries where expiration timestamp < current timestamp
- # For each selected dependency table entry:
  - < Delete HFS file identified by HFS filename for LOB
  - < Delete dependency table entry

## 7. More Advanced Cache Management

The management system provides a macro that allows a system administrator to execute a management system provided stored procedure to perform cache management. Use of this macro is referred to as more advanced cache management.

5 When more advanced cache management occurs, all cached dynamic web pages and LOBs with expiration timestamps that precede the current timestamp are automatically deleted.

The management system-provided macro for the more advanced management of the caches permits the user to perform the following types of management. First, a cleanup operation may be requested that deletes all cached dynamic web pages and LOBs with expiration timestamps that precede the current timestamp. Second, web pages may be deleted from the dynamic web page cache based on system administrator provided input, such as macro file and block name template values, timestamp values or a combination of name and timestamp values (e.g., delete web pages where key of cached page is like X and creation timestamp is less than some provided timestamp value Y). All cached LOBs that are referenced by the deleted web pages are also deleted. Third, LOBs from the LOB cache may be deleted based on a timestamp value (e.g., delete LOBs where creation timestamp is greater than some provided timestamp value Y). All cached dynamic web pages that contain hypertext references to the LOBs are first deleted.

## 8. Macro for More Advanced Cache Management

More specifically, the macro for more advanced management of dynamic web pages and LOBs allows a system administrator to specify the following:

- # The deletion of all cached dynamic web pages and LOBs with expiration timestamps that precede the current timestamp.
- 25 # A macro and block template value, where all dynamic web pages are deleted where ACTUAL\_KEY LIKE %value%.
- # A starting creation timestamp value: where all dynamic web pages are deleted where creation timestamp > or equal to starting creation timestamp.
- # An ending creation timestamp value: where all dynamic web pages are deleted

0980  
0240  
0240  
0240  
0240  
0240  
0240  
15  
20

5

- where creation timestamp < or equal to ending creation timestamp.
- # Both a starting and an ending creation timestamp value, where all dynamic web pages are deleted where creation timestamp > or equal to starting creation timestamp and < or equal to ending creation timestamp.
- # A macro and block template value and a starting creation timestamp value, where all dynamic web pages are deleted where ACTUAL\_KEY LIKE %value% and where creation timestamp > or equal to starting creation timestamp.
- # A macro and block template value and an ending creation timestamp value where all dynamic web pages are deleted where ACTUAL\_KEY LIKE %value% and where creation timestamp < or equal to ending creation timestamp.
- # A macro and block template value and both a starting and an ending creation timestamp value, where all dynamic web pages are deleted where ACTUAL\_KEY LIKE %value% and where creation timestamp > or equal to starting creation timestamp and < or equal to ending creation timestamp.

The macro for more advanced management of LOBs also allows a user to specify the following:

- # A starting creation timestamp value, where all LOBs are deleted where creation timestamp > or equal to starting creation timestamp.
- # An ending creation timestamp value, where all LOBs are deleted where creation timestamp < or equal to ending creation timestamp.
- # Both a starting and an ending creation timestamp value, where all LOBs are deleted where creation timestamp > or equal to starting creation timestamp and < or equal to ending creation timestamp.

The macro for more advanced management is named manage\_cache.d2w.

25 Example 1: To delete selected dynamic Web pages and related large objects that have been cached:  
a. Invoke the BEGIN HTML block of the manage\_cache.d2w macro.

b. Click on the Delete selected dynamic Web pages and related large object files choice.

This choice lets you specify a filter and timestamp values for the cached Web pages you want to delete. All expired cached Web pages and LOBs as well as all LOBs referenced by the deleted Web pages are also deleted.

5 c. Optionally type a string in the Enter the ACTUAL\_KEY filter field that matches any part of the ACTUAL\_KEY for the Web pages to be deleted. This string acts as filter for selecting the cached Web pages Net.Data deletes. The string can contain up to 250 characters.

For example, when the following string is entered:

/netdata/macros/my\_macro.d2w/report

Net.Data deletes all cached Web pages that have an ACTUAL\_KEY value containing this string.

d. Optionally click on the Starting CREATION\_TIME check box and enter a timestamp value. Net.Data deletes all cached Web pages that have creation times greater than or equal to this timestamp value, and that have creation times less than or equal to the Ending CREATION\_TIME, if specified. If no Ending CREATION\_TIME value is specified, than Net.Data deletes all cached Web pages that have creation times greater than or equal to the Starting CREATION\_TIME value. For example, when the following timestamp is entered:

20 Starting CREATION\_TIME:

Year 1999 Month 03 Day 23 Hour 14 Minute 00 Second 00

Net.Data deletes all cached Web pages that were created on or after 2:00 PM on March 23, 1999 up to and including the value of Ending CREATION\_TIME, if specified.

5 e. Optionally click on the Ending CREATION\_TIME check box and enter a timestamp value. Net.Data deletes all cached Web pages that have creation times less than or equal to this timestamp value, and that have creation times greater than or equal to the Starting CREATION\_TIME, if specified. If no Starting CREATION\_TIME value is specified, then Net.Data deletes all cached Web pages that have creation times less than or equal to the Ending CREATION\_TIME value. For example, when the following timestamp is entered:

Ending CREATION\_TIME:

Year 1999 Month 03 Day 23 Hour 23 Minute 59 Second 59

610 Net.Data deletes all cached Web pages that were created on or before 11:59:59 PM on March 23, 1999, starting with the value of Starting CREATION\_TIME, if specified.

615 If the Enter the ACTUAL\_KEY filter field is empty and neither of the check boxes are checked, Net.Data deletes only expired cached Web pages and LOBs.

620 f. Click the EXECUTE push button to proceed, or select Back to the beginning to return to the main page and cancel your request.

625 Example 2: To delete large objects in HFS and related Web pages that have been cached:

(i) Invoke the BEGIN HTML block of the manage\_cache.dtw macro.

(ii) Click on the Delete selected large object files and related dynamic Web pages choice.

70 This choice lets you specify timestamp values for the LOBs you want to delete.

All expired cached Web pages and LOBs as well as all LOBs referenced by deleted Web pages are also deleted.

75 (iii) Optionally click on the Starting CREATION\_TIME check box and enter a timestamp value. Net.Data deletes all LOB files that have creation times greater than or equal to this timestamp value, and that have creation times less than or equal to the Ending CREATION\_TIME, if specified. If no Ending

CREATION\_TIME value is specified, then Net.Data deletes all LOB files that have creation times greater than or equal to the Starting CREATION\_TIME value. For example, when the following timestamp is entered:

## Starting CREATION\_TIME

5 Year 1999 Month 03 Day 23 Hour 14 Minute 00 Second 00

Net.Data deletes all LOBs that were created on or after 2:00 PM on March 23, 1999, up to and including the value of Ending CREATION TIME, if specified.

(iv) Optionally click on the Ending CREATION\_TIME check box and enter a timestamp value. Net.Data deletes all LOB files that have creation times less than or equal to this timestamp value, and that have creation times greater than or equal to the Starting CREATION\_TIME, if specified. If no Starting CREATION\_TIME value is specified, then Net. Data deletes all LOB files that have creation times less than or equal to the Ending CREATION\_TIME value. For example, when the following timestamp is entered:

### Ending CREATION TIME:

Year 1999 Month 03 Day 23 Hour 23 Minute 59 Second 59

Net.Data deletes all LOB files that were created on or before 11:59:59 PM on March 23, 1999, starting with the value of Starting CREATION TIME, if specified.

If neither of the check boxes are checked, Net.Data deletes only expired LOB files and cached Web pages.

(v) Click the EXECUTE push button to proceed, or select Back to the

beginning to return to the main page and cancel your request.

## 9. Security Considerations and Guidelines

The userid(s) associated with the threads or processes that execute management system requests must have the EXECUTE privilege for the stored procedure used to add dynamic web pages to the cache and to add dependency information to the dynamic web page/LOB dependency table. Additionally, the userid(s) associated with the threads or processes that execute management system requests must have the authority to add files to the HTML\_PATH directory.

The userid(s) associated with the threads or processes that execute the management system requests must also have the EXECUTE privilege for the stored procedure used to automatically manage the dynamic web page and LOB caches. The userid(s) used for more advanced management of LOBs and cached web pages must have the EXECUTE privilege for the stored procedure used to manually manage the dynamic web page and LOB caches.

An administrative userid is typically used to create the HTML\_PATH directory and prepare the stored procedures for execution. The administrative userid must have the INSERT, SELECT, and DELETE privileges on the web cache table and on the dynamic web page/LOB dependency table.

In summary, the preferred embodiment of the present invention provides a technique for managing the contents of a dynamic web page system cache and related data stores containing objects that the web pages reference while also ensuring the completeness of cached web pages that are subsequently reused for display at browsers. The content may be managed by expiration times or in more advanced ways. Fig. 2 is a flow chart illustrating the steps performed in accordance with an embodiment of the present invention.

Block 200 represents the present invention receiving a request to generate a dynamic web page. The present invention then retrieves data and places the data in a dynamically generated web page, as represented by Block 202. The retrieved data may be linked to other stored data. Block 204 represents the present invention caching the retrieved data and the linked data. The present invention then manages the cached data, as represented by block 206.

## CONCLUSION

This concludes the description of the preferred embodiment of the invention. The following describes some alternative embodiments for accomplishing the present invention. For example, any type of computer, such as a mainframe, minicomputer, or personal computer, or 5 computer configuration, such as a timesharing mainframe, local area network, or standalone personal computer, could be used with the present invention.

The foregoing description of the preferred embodiment of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible 10 in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims appended hereto.

05000242-022380